arXiv:1511.05604v1 [stat.CO] 17 Nov 2015

blavaan: Bayesian Structural Equation Models via Parameter Expansion

Edgar C. Merkle
University of Missouri

Yves Rosseel
Ghent University

Abstract

This article describes a novel parameter expansion approach to Bayesian structural equation model (SEM) estimation in JAGS, along with an R package implementing the approach in an open fashion. The methodology and software are intended to provide users with a general means of estimating Bayesian SEMs, both classical and novel, in a straightforward fashion. Users can estimate Bayesian versions of classical SEMs using **lavaan** syntax, they can obtain state-of-the-art Bayesian fit measures associated with the models, and they can export JAGS code to modify the classical SEMs as desired. This paper explains and illustrates the parameter expansion approach, and it also provides an overview and examples of package **blavaan**.

Keywords: Bayesian SEM, structural equation models, JAGS, MCMC, **lavaan**.

The intent of **blavaan** is to implement Bayesian structural equation models (SEMs) that are satisfactory on all three of the following dimensions: (i) the speed with which models can be estimated, (ii) the ease by which models can be specified, and (iii) the ease by which the models can be extended to new situations. Existing software of which we are aware is satisfactory on a maximum of two of these dimensions. Software like Mplus (Muthén and Muthén 1998–2012) or AMOS (Arbuckle 2013) allows for easy specification and very fast estimation of Bayesian SEMs, but they do not allow the user to extend the models to novel situations (beyond what the software authors offer). Software like BUGS (Lunn, Jackson, Best, Thomas, and Spiegelhalter 2012), JAGS (Plummer 2003), and Stan (Stan Development Team 2014) allow for many novel model extensions, but the model specification can be tedious and the estimation can be slow. Finally, (2007; see also Song and Lee 2012) provides WinBUGS code for estimating many Bayesian structural equation models using conjugate priors, where the posteriors all have simple forms. This results in fast model estimation, but it cannot be used when we have non-diagonal residual covariance matrices (associated with observed variables). Thus, the method cannot be applied to all structural equation models.

In this paper, we describe an R package that allows for easier specification of Bayesian SEMs while also allowing for extensions as desired and yielding models that can be estimated relatively quickly (i.e., as quickly as JAGS would allow). As further described below, the approach builds on the work of many previous researchers: the model specification syntax of R package **lavaan** (Rosseel 2012); the approach of Lee (2007) for estimating models in WinBUGS; the approach of Palomo, Dunson, and Bollen (2007) for handling correlated residuals; and the approaches of Barnard, McCulloch, and Meng (2000) and Muthén and Asparouhov (2012) for specifying prior distributions of covariance matrices.

The resulting package, **blavaan**, is potentially useful for the analyst that wishes to estimate Bayesian SEMs and for the methodologist that wishes to develop novel extensions of Bayesian SEMs. The package is essentially a series of bridges between several existing R packages, with the bridges being used to make each step of the modeling easy and flexible. Package **lavaan** serves as the starting point and ending point in the sequence: the user specifies models via **lavaan** syntax, and objects of class **blavaan** make use of many **lavaan** functions. Consequently, writings on **lavaan** (e.g., Beaujean 2014; Rosseel 2012) give the user a good idea about how to use **blavaan**.

Following model specification, **blavaan** examines the details of the specified model and converts the specification to JAGS syntax. Importantly, this conversion makes use of the parameter expansion ideas presented below, resulting in MCMC chains that quickly converge for many models. We employ package **runjags** (Denwood in press) for sampling parameters and summarizing the MCMC chains. Once **runjags** is finished, **blavaan** organizes summaries and computes a variety of Bayesian fit measures. Finally, these results are returned in a **blavaan** object that can make use of many **lavaan** functions.

1. Structural Equation Model

We assume an $n \times p$ data matrix $\boldsymbol{Y} = (\boldsymbol{y}_1 \ \boldsymbol{y}_2 \ \ldots \ \boldsymbol{y}_n)'$ with n independent cases and p continuous manifest variables. A structural equation model with m latent variables may be represented by the equations

$$\boldsymbol{y} = \boldsymbol{\nu} + \boldsymbol{\Lambda}\boldsymbol{\eta} + \boldsymbol{\epsilon} \tag{1}$$

$$\boldsymbol{\eta} = \boldsymbol{\alpha} + \boldsymbol{B}\boldsymbol{\eta} + \boldsymbol{\zeta}, \tag{2}$$

where $\boldsymbol{\eta}$ is an $m \times 1$ vector containing the latent variables; $\boldsymbol{\epsilon}$ is a $p \times 1$ vector of residuals; and $\boldsymbol{\zeta}$ is an $m \times 1$ vector of residuals associated with the latent variables. Each entry in these three vectors is independent of the others. Additionally, $\boldsymbol{\nu}$ and $\boldsymbol{\alpha}$ contain intercept parameters for the manifest and latent variables, respectively; $\boldsymbol{\Lambda}$ is a matrix of factor loadings; and $\boldsymbol{B}$ contains parameters that reflect directed paths between latent variables.

In conventional SEMs, we assume multivariate normality of the $\boldsymbol{\epsilon}$ and $\boldsymbol{\zeta}$ vectors. In particular,

$$\boldsymbol{\epsilon} \sim N_p(\boldsymbol{0}, \boldsymbol{\Theta}) \tag{3}$$

$$\boldsymbol{\zeta} \sim N_m(\boldsymbol{0}, \boldsymbol{\Psi}), \tag{4}$$

with the latter equation implying multivariate normality of the latent variables. This framework can be extended to situations where some of the manifest variables are discrete (e.g., Song and Lee 2001, 2012), and we plan to include this functionality in future package versions. However, in this paper, we focus on models for continuous $\boldsymbol{y}$.

2. Bayesian SEM

As noted above, the approach of Lee (2007; see also Song and Lee 2012) involves the use of conjugate priors on the parameters from Equations (1) to (4). We provide an overview of the prior distributions below, though our description differs slightly from that of the original

authors. In particular, we do not distinguish between exogenous and endogenous latent variables.

Lee (2007) assumes that $\boldsymbol{\Theta}$ is diagonal and takes

$$\theta_{kk} \sim IG(a_k, b_k) \text{ for } k = 1, \ldots, p,$$

where $IG()$ is the inverse gamma distribution and a_k and b_k are prior parameters set by the analyst. He then sets multivariate normal priors on the parameters in $\boldsymbol{\nu}$ and on the parameters in $\boldsymbol{\Lambda}$, $\boldsymbol{\alpha}$, and $\boldsymbol{B}$. For the latter three matrices, Lee blocks the parameters by manifest variable and then places a multivariate normal prior on each block of parameters (while keeping track of fixed elements in each matrix). The prior covariance matrix of each block is scaled by the corresponding diagonal parameter of $\boldsymbol{\Theta}$.

Finally, Lee separately handles two parts of the latent variable covariance matrix $\boldsymbol{\Psi}$, the first part corresponding to exogenous variables and the second part corresponding to endogenous variables. The covariance matrix of exogenous latent variables is assigned an inverse Wishart prior, whereas the covariance matrix of endogenous latent variables is assumed to be diagonal with inverse gamma priors on each element. Covariances between exogenous and endogenous latent variables are assumed to be zero, completing the $\boldsymbol{\Psi}$ matrix.

While the above approach results in posterior distributions with simple forms (normal, inverse gamma, and inverse Wishart), the assumption of diagonal covariance matrices is restrictive in some applications. For example, researchers often wish to fit models with correlated residuals, and it has been argued that such models are both necessary and under-utilized (Cole, Ciesla, and Steiger 2007). Correlated residuals pose a difficult problem for MCMC methods because they often result in covariance matrices with some (but not all) off-diagonal entries equal to zero. In this situation, we cannot assign an inverse Wishart prior to the full covariance matrix because this does not allow us to fix some off-diagonal entries to zero. Further, if we assign a prior to each free entry of the covariance matrix and sample the parameters separately, we can often obtain covariance matrices that are not positive definite. This can lead to software crashes.

To address the issue of non-diagonal covariance matrices, Muthén and Asparouhov (2012) implemented a random walk method that is based on work by Chib and Greenberg (1998). This method samples free parameters of the covariance matrix via Metropolis-Hastings steps. While the implementation is fast and efficient, it does not allow for some types of equality constraints because parameters are updated in blocks (either all parameters in a block must be constrained to equality, or no parameters in a block can be constrained). Further, the method is unreliable for models involving many latent variables. Consequently, Muthén and Asparouhov (2012) implemented three other MCMC methods that are suited to different types of models.

In our initial work on package **blavaan**, we developed methodology for fitting models with free residual covariance parameters. We sought methodology that (i) would work in JAGS, (ii) was reasonably fast, and (iii) allowed for satisfactory specification of prior distributions. In the next section, we describe the resulting methodology. We also describe situations where we can avoid parameter expansion, which speeds up model estimation. It is important to note that, when we use terms like "reasonably fast," we mean that the estimation time is on the order of minutes (often, less than two minutes on desktop computers). This will never beat compiled code, but it is often faster than alternative JAGS parameterizations that rely on multivariate distributions.

2.1. Parameter Expansion

General parameter expansion approaches to Bayesian inference are described by Gelman (2004, 2006), with applications to factor analysis models detailed by Ghosh and Dunson (2009). Our approach here is related to that of Palomo *et al.* (2007), who employ phantom latent variables (they use the term pseudo-latent variable) to simplify the estimation of models with non-diagonal $\boldsymbol{\Theta}$ matrices. This results in a *working model* that is estimated, with the working model parameters being transformed to the *inferential model* parameters from Equations (1) and (2).

Overview

Assuming v nonzero entries in the lower triangle of $\boldsymbol{\Theta}$ (i.e., v residual covariances), Palomo *et al.* (2007) take the inferential model residual vector $\boldsymbol{\epsilon}$ and reparameterize it as

$$\begin{aligned}\boldsymbol{\epsilon} &= \boldsymbol{\Lambda}_D \boldsymbol{D} + \boldsymbol{\epsilon}^* \\ \boldsymbol{D} &\sim N_v(\mathbf{0}, \boldsymbol{\Psi}_D) \\ \boldsymbol{\epsilon}^* &\sim N_p(\mathbf{0}, \boldsymbol{\Theta}^*),\end{aligned}$$

where $\boldsymbol{\Lambda}_D$ is a $p \times v$ matrix with many zero entries, $\boldsymbol{D}$ is a $v \times 1$ vector of phantom latent variables, and $\boldsymbol{\epsilon}^*$ is a $p \times 1$ residual vector. This approach is useful because, by carefully choosing the nonzero entries of $\boldsymbol{\Lambda}_D$, both $\boldsymbol{\Psi}_D$ and $\boldsymbol{\Theta}^*$ are diagonal. This allows us to employ an approach related to Lee (2007), avoiding high-dimensional normal distributions in favor of univariate normals.

Under this working model parameterization, the inferential model covariance matrix $\boldsymbol{\Theta}$ can be re-obtained via

$$\boldsymbol{\Theta} = \boldsymbol{\Lambda}_D \boldsymbol{\Psi}_D \boldsymbol{\Lambda}_D' + \boldsymbol{\Theta}^*.$$

We can also handle covariances between latent variables in the same way, taking

$$\begin{aligned}\boldsymbol{\zeta} &= \boldsymbol{B}_E \boldsymbol{E} + \boldsymbol{\zeta}^* \\ \boldsymbol{E} &\sim N_w(\mathbf{0}, \boldsymbol{\Psi}_E) \\ \boldsymbol{\zeta}^* &\sim N_m(\mathbf{0}, \boldsymbol{\Psi}^*),\end{aligned}$$

where the original covariance matrix $\boldsymbol{\Psi}$ is re-obtained in a similar fashion:

$$\boldsymbol{\Psi} = \boldsymbol{B}_E \boldsymbol{\Psi}_E \boldsymbol{B}_E' + \boldsymbol{\Psi}^*.$$

This approach to handling latent variable covariances tends to be suboptimal in models with many (say, > 5) correlated latent variables. In these cases, it is better to sample the latent variables from multivariate normal distributions. Package **blavaan** attempts to use multivariate normal distributions in situations where it can, reverting to the above reparameterization only when necessary.

To choose nonzero entries of $\boldsymbol{\Lambda}_D$ (with the same method applying to nonzero entries of $\boldsymbol{B}_E$), we define two $v \times 1$ vectors $\boldsymbol{r}$ and $\boldsymbol{c}$. These vectors contain the respective row numbers and column numbers of the nonzero, lower-triangular entries of $\boldsymbol{\Theta}$. For $j = 1, \ldots, p$, the nonzero entries of $\boldsymbol{\Lambda}_D$ then occur in column j, rows r_j and c_j. Palomo *et al.* (2007) set the nonzero entries equal to 1, which is suboptimal because it does not allow for negative covariances.

Instead of 1s, we put extra parameters in $\mathbf{\Lambda}_D$ that result in well-defined priors on the original covariance matrix $\mathbf{\Theta}$. This is further described in the next section. First, however, we give an example of the approach.

Example

Consider the famous political democracy data from Bollen (1989), which includes eleven observed variables that are hypothesized to arise from three latent variables. The inferential structural equation model affiliated with these data appears as a path diagram (with variance parameters omitted) at the top of Figure 1. This model can be written in matrix form as

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \\ y_9 \\ y_{10} \\ y_{11} \end{pmatrix} = \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \\ \nu_5 \\ \nu_6 \\ \nu_7 \\ \nu_8 \\ \nu_9 \\ \nu_{10} \\ \nu_{11} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ \lambda_2 & 0 & 0 \\ \lambda_3 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \lambda_5 & 0 \\ 0 & \lambda_6 & 0 \\ 0 & \lambda_7 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & \lambda_5 \\ 0 & 0 & \lambda_6 \\ 0 & 0 & \lambda_7 \end{pmatrix} \begin{pmatrix} \text{ind60} \\ \text{dem60} \\ \text{dem65} \end{pmatrix} + \boldsymbol{\epsilon}$$

$$\begin{pmatrix} \text{ind60} \\ \text{dem60} \\ \text{dem65} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ b_1 & 0 & 0 \\ b_2 & b_3 & 0 \end{pmatrix} \begin{pmatrix} \text{ind60} \\ \text{dem60} \\ \text{dem65} \end{pmatrix} + \boldsymbol{\zeta}.$$

Additionally, the latent residual covariance matrix $\mathbf{\Psi}$ is diagonal and the observed residual covariance matrix is

$$\mathbf{\Theta} = \begin{pmatrix} \theta_{11} & & & & & & & & & & \\ & \theta_{22} & & & & & & & & & \\ & & \theta_{33} & & & & & & & & \\ & & & \theta_{44} & & & & \theta_{48} & & & \\ & & & & \theta_{55} & & \theta_{57} & & \theta_{59} & & \\ & & & & & \theta_{66} & & & & \theta_{6,10} & \\ & & & & \theta_{57} & & \theta_{77} & & & & \theta_{7,11} \\ & & & \theta_{48} & & & & \theta_{88} & & & \\ & & & & \theta_{59} & & & & \theta_{99} & & \theta_{9,11} \\ & & & & & \theta_{6,10} & & & & \theta_{10,10} & \\ & & & & & & \theta_{7,11} & & \theta_{9,11} & & \theta_{11,11} \end{pmatrix},$$

where the missing entries all equal zero. The six unique, off-diagonal entries of $\mathbf{\Theta}$ make the model difficult to estimate via Bayesian methods.

To ease model estimation, we follow Palomo *et al.* (2007) in reparameterizing the model to the working model displayed at the bottom of Figure 1. The working model is now

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \\ y_9 \\ y_{10} \\ y_{11} \end{pmatrix} = \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \\ \nu_5 \\ \nu_6 \\ \nu_7 \\ \nu_8 \\ \nu_9 \\ \nu_{10} \\ \nu_{11} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ \lambda_2 & 0 & 0 \\ \lambda_3 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \lambda_5 & 0 \\ 0 & \lambda_6 & 0 \\ 0 & \lambda_7 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & \lambda_5 \\ 0 & 0 & \lambda_6 \\ 0 & 0 & \lambda_7 \end{pmatrix} \begin{pmatrix} \text{ind60} \\ \text{dem60} \\ \text{dem65} \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda_{D1} & 0 & 0 & 0 & 0 & 0 \\ 0 & \lambda_{D2} & \lambda_{D3} & 0 & 0 & 0 \\ 0 & 0 & 0 & \lambda_{D4} & 0 & 0 \\ 0 & \lambda_{D5} & 0 & 0 & \lambda_{D6} & 0 \\ \lambda_{D7} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \lambda_{D8} & 0 & 0 & \lambda_{D9} \\ 0 & 0 & 0 & \lambda_{D10} & 0 & 0 \\ 0 & 0 & 0 & 0 & \lambda_{D11} & \lambda_{D12} \end{pmatrix} \begin{pmatrix} D1 \\ D2 \\ D3 \\ D4 \\ D5 \\ D6 \end{pmatrix} + \boldsymbol{\epsilon}^*$$

$$\begin{pmatrix} \text{ind60} \\ \text{dem60} \\ \text{dem65} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ b_1 & 0 & 0 \\ b_2 & b_3 & 0 \end{pmatrix} \begin{pmatrix} \text{ind60} \\ \text{dem60} \\ \text{dem65} \end{pmatrix} + \boldsymbol{\zeta}.$$

The covariance matrix associated with $\boldsymbol{\epsilon}^*$, $\boldsymbol{\Theta}^*$, is now diagonal, which makes the model easier to estimate via MCMC. The difference between our approach and that of Palomo *et al.* (2007) is that we estimate the loadings involving the "D" latent variables, whereas Palomo *et al.* (2007) fix these loadings to 1. This allows us to obtain negative estimates of residual covariances if needed, which is an improvement over the previous approach. Estimation of these loadings comes with a cost, however, in that the prior distributions of the residual covariance parameters become difficult to meaningfully set. This problem and a solution are further described in the next section.

2.2. Priors on Covariances

Due to the reparameterization of covariances, prior distributions on the working model parameters require care in their specification. Every inferential model covariance parameter is potentially the product of three parameters: two representing paths from a phantom latent variable to the variable of interest, and one precision (variance) of the phantom latent variable. These three parameters also impact the inferential model variance parameters. We carefully restrict these parameters to arrive at prior distributions that are meaningful, that are flexible, and that allow for fast JAGS model estimation.

We place univariate prior distributions on model variance and correlation parameters, with the distributions being based on the inverse Wishart or on other prior distributions described in the literature. This is highly related to the approach of Barnard *et al.* (2000), who separately specify prior distributions on the variance and correlation parameters comprising a covariance matrix. Our approach is also related to Muthén and Asparouhov (2012), who focus on marginal distributions associated with the inverse Wishart. A notable difference is that we always use proper prior distributions, whereas Muthén and Asparouhov (2012) often use improper prior distributions.

To illustrate the approach, we refer to Figure 2. This figure uses path diagrams to display two observed variables with correlated residuals (left panel) along with their parameter-expanded representation in the working model (right panel). Our goal is to specify a sensible prior distribution for the inferential model's variance/covariance parameters in the left panel, using the parameters in the right panel. The covariance of interest, θ_{12}, has been converted to a normal phantom latent variable with variance ψ_D and two directed paths, λ_1 and λ_2. To

Figure 1: Political democracy path diagrams, original (top) and parameter-expanded version (bottom). Variance parameters are omitted.

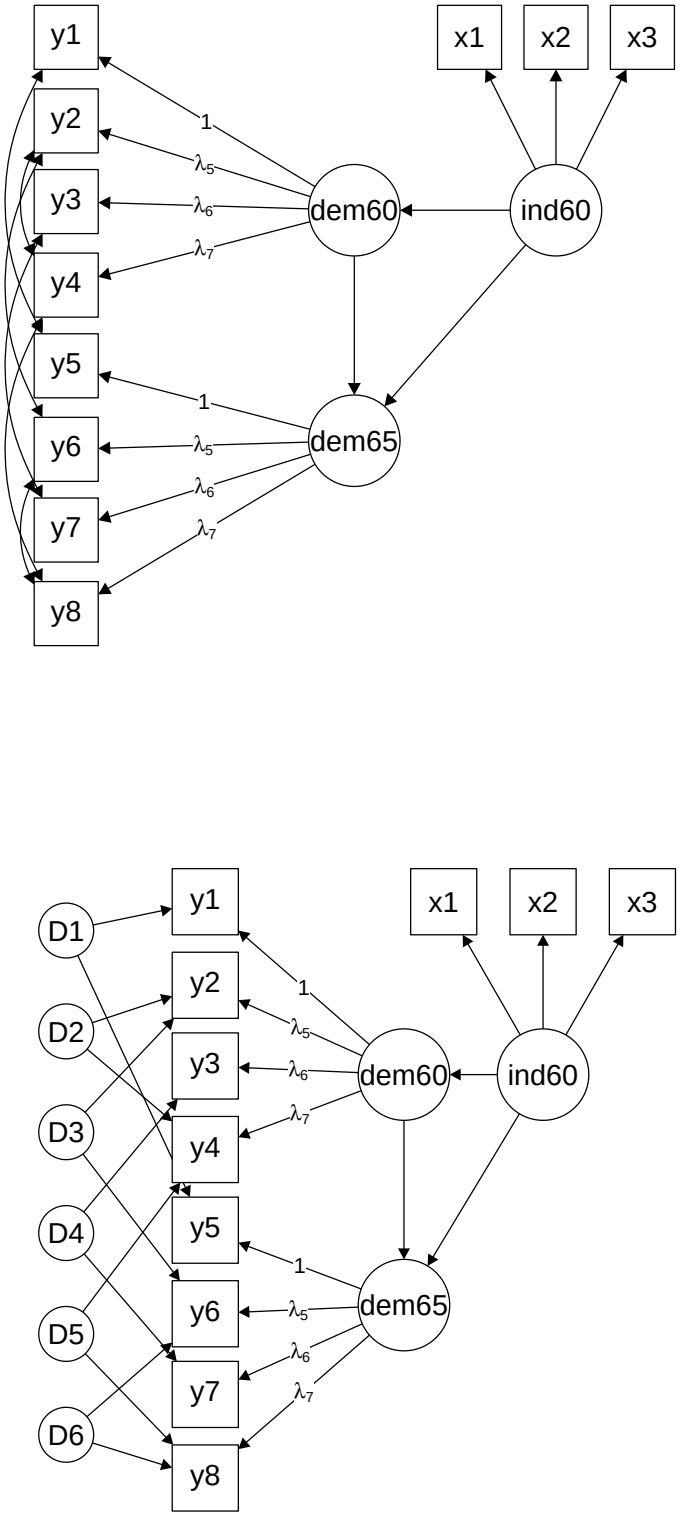

implement an approach related to that of Barnard *et al.* (2000) in JAGS, we set

$$\begin{aligned}\psi_D &= 1\\ \lambda_1 &= \sqrt{|\rho_{12}|\theta_{11}}\\ \lambda_2 &= \text{sign}(\rho_{12})\sqrt{|\rho_{12}|\theta_{22}}\\ \theta^*_{11} &= \theta_{11} - |\rho_{12}|\theta_{11}\\ \theta^*_{22} &= \theta_{22} - |\rho_{12}|\theta_{22},\end{aligned}$$

where ρ_{12} is the correlation associated with the covariance θ_{12}, and sign(a) equals 1 if $a > 0$ and -1 otherwise.

Using this combination of parameters, we need only set prior distributions on the inferential model parameters θ_{11}, θ_{22}, and ρ_{12}, with θ_{12} being obtained from these parameters in the usual way. If we want our priors to be analogous to an inverse Wishart prior with d degrees of freedom and diagonal scale matrix $\boldsymbol{S}$, we can set univariate priors of

$$\begin{aligned}\theta_{11} &\sim IG((d-p+1)/2, s_{11}/2)\\ \theta_{22} &\sim IG((d-p+1)/2, s_{22}/2)\\ \rho_{12} &\sim \text{Beta}_{(-1,1)}((d-p+1)/2, (d-p+1)/2),\end{aligned}$$

where $\text{Beta}_{(-1,1)}$ is a beta distribution with support on $(-1, 1)$ instead of the usual $(0, 1)$ and p is the dimension of $\boldsymbol{S}$ (typically, the number of observed variables). These priors are related to those used by Muthén and Asparouhov (2012), based on results from Barnard *et al.* (2000). They are also the default priors for variance/correlation parameters in **blavaan**, with $d = (p + 1)$ and $\boldsymbol{S} = \boldsymbol{I}$. We refer to this parameterization option as `srs`, reflecting the fact that we are dissecting the covariance parameters into standard deviation and correlation parameters.

Barnard *et al.* (2000) avoid the inverse gamma priors on the variance parameters, instead opting for log-normal distributions on the standard deviations that they judged to be more numerically stable. Package **blavaan** allows for custom priors, so that the user can employ the log-normal or others on precisions, variances, or standard deviations (this is illustrated later). This approach is generally advantageous because it allows for flexible specification of prior distributions on covariance matrices, including those with fixed zeros or those where we have different amounts of prior information on certain variance/correlation parameters within the matrix.

While we view the `srs` approach as optimal for dealing with covariance parameters here, it is not the fastest possible approach in JAGS. It can sometimes be faster to treat the phantom latent variables similarly to the other latent variables in the model, assigning prior distributions in the same manner. This approach, which we call `fa` (because the priors resemble traditional factor analysis priors), involves the following default prior distributions on the working model from Figure 2:

$$\begin{aligned}\psi_D &\sim IG(1, .5)\\ \lambda_1 &\sim N(0, 10^4)\\ \lambda_2 &\sim N(0, 10^4)\\ \theta^*_{11} &\sim IG(1, .5)\\ \theta^*_{22} &\sim IG(1, .5),\end{aligned}$$

Figure 2: Example of phantom latent variable approach to covariances.

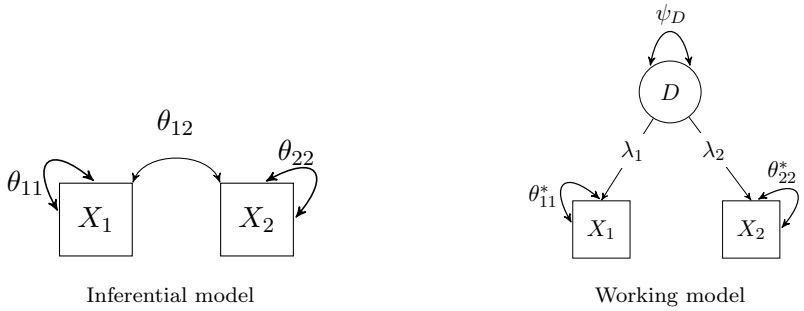

with the inferential model parameters being obtained in the usual way:

$$\theta_{11} = \theta^*_{11} + \lambda_1^2 \psi_D$$
$$\theta_{22} = \theta^*_{22} + \lambda_2^2 \psi_D$$
$$\theta_{12} = \lambda_1 \lambda_2 \psi_D.$$

This approach can be faster in JAGS because the working model prior distributions are all of the conjugate type, which can avoid the need for model adaptation. The main disadvantage of the approach is that the implied prior distributions on the inferential model parameters are difficult to describe. Thus, it is difficult to introduce informative prior distributions for the inferential model variance/covariance parameters. Additionally, the working model parameters are not identified by the likelihood, which can complicate Laplace approximation of the marginal likelihood (which, as described below, is related to the Bayes factor). This is because the approximation works best when the posterior distributions are unimodal.

In **blavaan**, the user is free to specify `srs` or `fa` priors separately for covariances between observed variables (argument `ov.cp`) and for covariances between latent variables (argument `lv.cp`). Additionally, the package attempts to identify when it can sample latent variables directly from a multivariate normal distribution. This most commonly happens when we have many exogenous latent variables that all covary with one another. In this case, we place an inverse-Wishart prior distribution on the latent variable covariance matrix by default for speed and efficiency.

While we prefer the `srs` approach to covariances between observed variables, the `fa` approach may be useful during initial modeling, when users are more interested in quickly obtaining rough results rather than obtaining the best results.

2.3. Model Fit & Comparison

Package **blavaan** supplies a variety of statistics for model evaluation and comparison, available via the `fitMeasures()` function. For model evaluation, it supplies posterior predictive checks of the model's log-likelihood (e.g., Gelman, Carlin, Stern, and Rubin 2004). For model comparision, it supplies a variety of statistics including the Deviance Information Criterion (DIC; Spiegelhalter, Best, Carlin, and van der Linde 2002), the (Laplace-approximated) log-Bayes factor, the Widely Applicable Information Criterion (WAIC; Watanabe 2010), and the leave-one-out cross-validation statistic (LOO; e.g., Gelfand 1996). The latter two statistics are computed by R package **loo** (Vehtari, Gelman, and Gabry 2015b), using output from **blavaan**.

Calculation of the information criteria are somewhat complicated by the fact that, during JAGS model estimation, we condition on the latent variables $\boldsymbol{\eta}_i$, $i = 1, \ldots, n$. That is, our model utilizes likelihoods of the form $L(\boldsymbol{\vartheta}, \boldsymbol{\eta}_i | \boldsymbol{y}_i)$, where $\boldsymbol{\vartheta}$ is the vector of model parameters and $\boldsymbol{\eta}_i$ is a vector of latent variables associated with individual i. Because the latent variables are random effects, we must integrate them out to obtain $L(\boldsymbol{\vartheta}|\boldsymbol{Y})$, the likelihood by which model assessment statistics are typically computed. This is easy to do for the models considered here, because the integrated likelihood continues to be multivariate normal. However, we cannot rely on JAGS to automatically calculate the correct likelihood, so that **blavaan** calculates the likelihood separately after parameters have been sampled in JAGS.

Posterior Predictive Checks

As a measure of the model's absolute fit, **blavaan** computes a posterior predictive p-value that compares observed likelihood ratio test (χ^2) statistics to likelihood ratio test statistics generated from the model's posterior predictive distribution. The likelihood ratio test statistic is generally computed via

$$\text{LRT}(\boldsymbol{Y}, \boldsymbol{\vartheta}) = -2 \log L(\boldsymbol{\vartheta}|\boldsymbol{Y}) + 2 \log L(\boldsymbol{\vartheta}^{\text{sat}}|\boldsymbol{Y}),$$

where $\boldsymbol{\vartheta}^{\text{sat}}$ is a "saturated" parameter vector that perfectly matches the observed mean and covariance matrix.

For each posterior draw $\boldsymbol{\vartheta}^s$ ($s = 1, \ldots, S$), computation of the posterior predictive p-value includes four steps:

1. Compute the observed LRT statistic as $\text{LRT}(\boldsymbol{Y}, \boldsymbol{\vartheta}^s)$.
2. Generate artificial data $\boldsymbol{Y}^{\text{rep}}$ from the model, assuming parameter values equal to $\boldsymbol{\vartheta}^s$.
3. Compute the posterior predictive LRT statistic $\text{LRT}(\boldsymbol{Y}^{\text{rep}}, \boldsymbol{\vartheta}^s)$.
4. Record whether or not $\text{LRT}(\boldsymbol{Y}^{\text{rep}}, \boldsymbol{\vartheta}^s) > \text{LRT}(\boldsymbol{Y}, \boldsymbol{\vartheta}^s)$.

The posterior predictive p-value is then the proportion of the time (out of S draws) that the posterior predictive LRT statistic is larger than the observed LRT statistic. Values closer to 1 indicate a model that fits the observed data, while values closer to 0 indicate the opposite.

This measure is impractical when there are missing data because it is impractical to compute the saturated log-likelihood in the presence of missing data. We need to use, e.g., the EM algorithm to do this, and this needs to be done $(S+1)$ times (once for the saturated likelihood of the observed data and once for the saturated likelihood of each of the S artificial datasets). With complete data, we can more simply calculate the saturated log-likelihood using the sample mean vector and covariance matrix. The argument `test="none"` can be supplied to **blavaan** in order to bypass these slow calculations when there are missing data.

DIC

The DIC is given as

$$DIC = -2 \log L(\hat{\boldsymbol{\vartheta}}|\boldsymbol{Y}) + 2 \text{ efp}_{\text{dic}}, \tag{5}$$

where $\hat{\boldsymbol{\vartheta}}$ is a vector of posterior parameter means (or other measure of central tendency), $L(\hat{\boldsymbol{\vartheta}}|\boldsymbol{Y})$ is the model's likelihood at the posterior means, and efp_{dic} is the model's effective

number of parameters, calculated as

$$\text{efp}_{\text{dic}} = 2\left[\log L(\hat{\boldsymbol{\vartheta}}|\boldsymbol{Y}) - \overline{\log L(\boldsymbol{\vartheta}|\boldsymbol{Y})}\right]. \tag{6}$$

The latter term, $\overline{\log L(\boldsymbol{\vartheta}|\boldsymbol{Y})}$, is obtained by calculating the log-likelihood for each posterior sample and then averaging.

Many readers will be familiar with the automatic calculation of DIC within programs like BUGS and JAGS. As we described above, the automatic DIC is not what users typically desire because the likelihood conditions on the latent variables $\boldsymbol{\eta}_i$. This greatly increases the effective number of parameters and may result in poor inferences (for further discussion of this issue, see Millar 2009). As previously mentioned, **blavaan** avoids the automatic DIC computation in JAGS, calculating its own likelihoods after model parameters have been sampled.

WAIC and LOO

The WAIC and LOO are asymptotically equivalent measures that are advantageous over DIC while also being more difficult to compute than DIC. Many computational difficulties are overcome in developments by Vehtari, Gelman, and Gabry (2015a), with their methodology being implemented in package **loo** (Vehtari *et al.* 2015b). As input, **loo** requires casewise log-likelihoods associated with a set of posterior samples: $\log L(\boldsymbol{\vartheta}^s|\boldsymbol{y}_i), s = 1, \ldots, S; i = 1, \ldots, n$.

Like DIC, both of these measures seek to characterize a model's predictive accuracy (generalizability). The definition of WAIC looks similar to that of DIC:

$$\text{WAIC} = -2\text{lppd} + 2\ \text{efp}_{\text{waic}},$$

where the first term is related to log-likelihoods of observed data and the second term involves an effective number of parameters. Both terms now involve log-likelihoods associated with individual observations, however, which help us estimate the model's expected log pointwise predictive density (a measure of predictive accuracy).

The first term, the log pointwise predictive density of the observed data (lppd), is estimated via

$$\text{lppd} = \sum_{i=1}^{n} \log\left(\frac{1}{S}\sum_{s=1}^{S} f(\boldsymbol{y}_i|\boldsymbol{\vartheta}^s)\right), \tag{7}$$

where S is the number of posterior draws and $f(\boldsymbol{y}_i|\boldsymbol{\vartheta}^s)$ is the density of observation i wrt the parameters sampled at iteration s. The second term, the effective number of parameters, is estimated via

$$\text{efp}_{\text{waic}} = \sum_{i=1}^{n} \text{var}_s(\log f(\boldsymbol{y}_i|\boldsymbol{\vartheta})),$$

where we compute a separate variance for each observation i across the S posterior draws.

The LOO measure estimates the predictive density of each individual observation from a cross-validation standpoint; that is, the predictive density when we hold out one observation at a time and use the remaining observations to update the prior. We can use analytical results to estimate this quantity based on the full-data posterior, removing the need to re-estimate the posterior while sequentially holding out each observation. These results lead to the estimate

$$\text{LOO} = -2\sum_{i=1}^{n} \log\left(\frac{\sum_{s=1}^{S} w_i^s f(\boldsymbol{y}_i|\boldsymbol{\vartheta}^s)}{\sum_{s=1}^{S} w_i^s}\right), \tag{8}$$

where the w_i^s are importance sampling weights based on the relative magnitude of individual i's density function across the S posterior samples. Package **loo** smooths these weights via a generalized Pareto distribution, improving the estimate.

We can also estimate the effective number of parameters under the LOO measure by comparing LOO to the lppd that was used in the WAIC calculation. This gives us

$$\text{efp}_{\text{loo}} = \text{lppd} + \text{LOO}/2,$$

where these terms come from Equations (7) and (8), respectively. Division of the latter term by two (and addition, vs subtraction) offsets the multiplication by -2 that occurs in (8). For further detail on all these measures, see Vehtari *et al.* (2015a).

Bayes Factor

The Bayes factor between two models is defined as the ratio of the models' marginal likelihoods (e.g., Kass and Raftery 1995). Candidate model 1's marginal likelihood can be written as

$$f_1(\boldsymbol{Y}) = \int f(\boldsymbol{\vartheta}_1) f(\boldsymbol{Y}|\boldsymbol{\vartheta}_1) \partial \boldsymbol{\vartheta}_1, \tag{9}$$

where $\boldsymbol{\vartheta}_1$ is a vector containing candidate model 1's free parameters (excluding the latent variables $\boldsymbol{\eta}_i$) and $f()$ is a probability density function. The marginal likelihood of candidate model 2 may be written in the same manner. The Bayes factor is then

$$BF_{12} = \frac{f_1(\boldsymbol{Y})}{f_2(\boldsymbol{Y})},$$

with values greater than 1 favoring Model 1.

Because the integral from (9) is generally difficult to calculate, Lewis and Raftery (1997) describe a Laplace-Metropolis estimator of the Bayes factor (also see Raftery 1993). This estimator relies on the Laplace approximation to integrals that involve a natural exponent. Such integrals can be written as

$$\int \exp(h(\boldsymbol{u})) \partial \boldsymbol{u} \approx (2\pi)^{Q/2} |\boldsymbol{H}^*|^{1/2} \exp(h(\boldsymbol{u}^*)),$$

where $h()$ is a function of the vector $\boldsymbol{u}$, Q is the length of $\boldsymbol{u}$, $\boldsymbol{u}^*$ is the value of $\boldsymbol{u}$ that maximizes h, and $\boldsymbol{H}^*$ is the inverse of the information matrix evaluated at $\boldsymbol{u}^*$. As applied to marginal likelihoods, we take $h(\boldsymbol{\vartheta}) = \log(f(\boldsymbol{\vartheta}) f(\boldsymbol{Y}|\boldsymbol{\vartheta}))$ so that our solution to Equation (9) is

$$f_1(\boldsymbol{Y}) \approx (2\pi)^{Q/2} |\boldsymbol{H}^*|^{1/2} f(\boldsymbol{\vartheta}^*) f(\boldsymbol{Y}|\boldsymbol{\vartheta}^*), \tag{10}$$

where $\boldsymbol{\vartheta}^*$ is a vector of posterior means (or other posterior estimate of central tendency), obtained from the MCMC output. In package **blavaan**, we follow Lewis and Raftery (1997) and compute the logarithm of this approximation for numerical stability.

Now that we have defined the model assessment statistics in **blavaan**, we provide further details on how the package works. We then illustrate the package via example.

3. Overview of blavaan

Readers familiar with **lavaan** will naturally be familiar with **blavaan**. For example, the main functions are the same as the main **lavaan** functions, except that they start with the letter

'b'. For example, confirmatory factor analysis models are estimated via `bcfa()` and structural equation models are estimated via `bsem()`. These two functions call the more general `blavaan()` function with a specific arrangement of arguments. The **blavaan** model specification syntax is nearly the same as the **lavaan** model specification syntax, and many functions are used by both packages.

There are two new features in **blavaan** that are specific to Bayesian modeling: prior distribution specification and the export/use of JAGS syntax. We discuss these topics by example, making use of a simple "error-in-variables" regression model. The model predicts students' verbal reasoning test scores based on their numerical reasoning test scores, using data from Dutch high school students who participated in a "stereotype threat" study (see Wicherts, Dolan, and Hessen 2005). Code for model specification and estimation appears below.

```
data(StereotypeThreat, package="psychotools")

model <- ' verbal ~ numerical '

## regression with non-fixed x
fit <- bsem(model, data=StereotypeThreat, fixed.x=FALSE)
```

3.1. Prior Distribution Specification

Package **blavaan** defines a default prior distribution for each type of model parameter (see Equations (1) and (2)). These default priors can be viewed via

```
dpriors()

##              nu           alpha          lambda            beta          itheta
## "dnorm(0,1e-3)" "dnorm(0,1e-2)" "dnorm(0,1e-2)" "dnorm(0,1e-2)"  "dgamma(1,.5)"
##            ipsi             rho           ibpsi
##  "dgamma(1,.5)"    "dbeta(1,1)" "dwish(iden,3)"
```

which includes a separate default prior for eight types of parameters. Default prior distributions are placed on precisions instead of variances, so that the letter i in `itheta` and `ipsi` stands for "inverse." The `ibpsi` prior is used for blocks of variance parameters that are all allowed to covary with one another; this block prior improves sampling efficiency and is currently applied to variances involving exogenous latent variables.

For most parameters, the user is free to set priors using any type of distribution that is defined in JAGS. Changes to default priors can be supplied with the `dp` argument. Further, the modifiers `[sd]` and `[var]` can be used to put a prior on a standard deviation or variance parameter, instead of the corresponding precision. For example,

```
fit <- bsem(model, data=StereotypeThreat, fixed.x=FALSE,
            dp=dpriors(nu="dnorm(5,1)", itheta="dunif(0,20)[sd]"))
```

sets the default priors for (i) the manifest variables' intercepts to normal with mean 5 and precision 1, and (ii) the manifest variables' standard deviations to uniform with bounds of 0 and 20.

Priors associated with the `rho` and `ibpsi` parameters are less amenable to change. The default `rho` prior is a beta(1,1) distribution with support on $(-1, 1)$; beta distributions with other parameter values could be used, but this distribution must be beta for now. The default prior on blocks of precision parameters is the Wishart with identity scale matrix and degrees of freedom equal to the dimension of the scale matrix plus one. The user cannot currently make changes to this prior distribution via the `dp` argument. Changes can be made, however, by exporting the JAGS syntax and manually editing it. This is further described in the next section.

In addition to priors on classes of model parameters, the user may wish to set the prior of a specific model parameter. This is accomplished by using the `prior()` modifier within the model specification. For example,

```
model <- ' verbal ~ prior("dnorm(0,1)")*numerical
           numerical ~~ prior("dlnorm(0,1)[sd]")*numerical '
```

sets specific priors for the regression weight and for the standard deviation of `numerical`. Priors set in this manner (via the model syntax) take precedence over the default priors set via the `dp` argument. Additionally, if one specifies a covariance parameter in the model syntax, the `prior()` modifier is for the correlation associated with the covariance, as opposed to the covariance itself. Thus, the prior should be a beta distribution.

3.2. JAGS Syntax

Users may find that some desired features are not currently implemented in **blavaan**; such features may be related to, e.g., the use of specific types of priors or the handling of discrete variables. While we plan to implement more features in the future, we also allow the user to export the JAGS syntax so that they can do it themselves. We believe this will be especially useful for researchers who wish to develop new types of models. These researchers can specify a traditional model in **blavaan** that is related to the one that they want to implement, and they can then obtain the JAGS syntax and data for the traditional model. This JAGS syntax should ease implementation of the novel model.

To export the JAGS syntax, users can employ the `jagfile` argument.

```
fit <- bsem(model, data=StereotypeThreat, fixed.x=FALSE, jagfile=TRUE)
```

If set to `TRUE`, the syntax will be written to the `lavExport` folder; if a character string is supplied, **blavaan** will use the character string as the folder name.

When the syntax is exported, two files are written within the folder. The first, `sem.jag`, contains the model specification in **JAGS** syntax. The second, `semjags.rda`, is a list named `jagtrans` that contains the data and initial values in JAGS format, as well as labels associated with model parameters. These pieces can be used to run the model "manually" via, e.g., package **rjags** or **runjags**. The example below illustrates how we could estimate the exported model via **runjags**.

```
## load jags data/inits/labels
load("lavExport/semjags.rda")
## model estimation
fit <- run.jags("lavExport/sem.jag", monitor=jagtrans$coefvec$jlabel,
                data=jagtrans$data, inits=jagtrans$inits)
```

If the user modifies the JAGS file, then the `monitor` and `inits` arguments might require modification (or, for automatic initial values from JAGS, the user could omit the `inits` argument). The data should not require modification unless the user adds or removes observed variables from the model.

We should also note that the dimensions of the parameter vectors/matrices in the JAGS syntax do not match the dimensions that we might expect from Equations (1) and (2). Instead, the number of rows reflect the number of model parameters of a certain type (and within a certain group), and the number of columns reflect the number of groups. For example, if we have a 3-group model with 9 factor loadings per group, then the JAGS syntax defines a `lambda` matrix with 9 rows and 3 columns.

Now that we have seen some features that are novel to **blavaan**, we will further illustrate the package by example.

4. Applications

In the applications below, we first illustrate some general features involving prior distribution specification and model fit measures. We then illustrate the study of measurement invariance via Bayesian models, also showing how we can use a fitted model of class **lavaan** to immediately estimate a model in **blavaan**.

4.1. Political Democracy

We begin with the Bollen (1989) industrialization-democracy model discussed earlier in the paper. The model specification is identical to the specification that we would use in **lavaan** to estimate the model via classical methods:

```
model <- '
       # latent variable definitions
         ind60 =~ x1 + x2 + x3
         dem60 =~ y1 + a*y2 + b*y3 + c*y4
         dem65 =~ y5 + a*y6 + b*y7 + c*y8

       # regressions
         dem60 ~ ind60
         dem65 ~ ind60 + dem60

       # residual correlations
         y1 ~~ y5
         y2 ~~ y4 + y6
         y3 ~~ y7
         y4 ~~ y8
         y6 ~~ y8
'
```

though we could also use the `prior()` modifier to include priors within the model specification. Instead of doing this, we specify priors for classes of parameters within the `bsem()` command below. We first place $N(5, \sigma^{-2} = .01)$ priors on the observed-variable intercepts (ν). Next, following Barnard *et al.* (2000), we place log-normal priors on the residual standard deviations of both observed and latent variables. This is accomplished by adding `[sd]` to the end of

the `itheta` and `ipsi` priors. Barnard et al. stated that the log-normal priors result in better computational stability, as opposed to the use of gamma priors on the precision parameters. Finally, we set beta(3,3) priors on the correlations associated with the covariance parameters. This reflects the belief that the correlations will generally be closer to 0 than to −1 or 1. The result is the following command:

```
fit <- bsem(model, data=PoliticalDemocracy,
            dp=dpriors(nu="dnorm(5,1e-2)", itheta="dlnorm(1,.1)[sd]",
                       ipsi="dlnorm(1,.1)[sd]", rho="dbeta(3,3)"),
            jagcontrol=list(method="rjparallel"))
```

where the latter argument, `jagcontrol`, allows us to sample each chain in parallel for faster model estimation. Defaults for the number of adaptation, burnin, and drawn samples come from package **runjags**; those defaults are currently 1,000, 4,000, and 10,000, respectively. Users can modify any of these via the arguments `adapt`, `burnin`, and `sample`.

Once we have sampled for the desired number of iterations, we can make use of **lavaan** functions in order to summarize and study the fitted model. Users may primarily be interested in `summary()`, which organizes results in a manner similar to the classical **lavaan** models.

```
summary(fit)

## blavaan (0.1-1) results of 10000 samples after 5000 adapt+burnin iterations
##
##   Number of observations                            75
##
##   Number of missing patterns                         1
##
##   Statistic                                 MargLogLik         PPP
##   Value                                      -2542.457       0.544
##
## Parameter Estimates:
##
##   Information                                     MCMC
##   Standard Errors                                 MCMC
##
## Latent Variables:
##                    Post.Mean  Post.SD  HPD.025  HPD.975     PSRF             Prior
##   ind60 =~
##     x1                 1.000
##     x2                 2.235    0.144    1.969    2.524    1.003     dnorm(0,1e-2)
##     x3                 1.835    0.161    1.511    2.145    1.001     dnorm(0,1e-2)
##   dem60 =~
##     y1                 1.000
##     y2          (a)    1.215    0.149    0.935    1.518    1.008     dnorm(0,1e-2)
##     y3          (b)    1.192    0.128    0.943     1.45    1.016     dnorm(0,1e-2)
##     y4          (c)    1.271    0.132    1.003    1.525    1.009     dnorm(0,1e-2)
##   dem65 =~
##     y5                 1.000
##     y6          (a)    1.215    0.149    0.935    1.518    1.008
##     y7          (b)    1.192    0.128    0.943     1.45    1.016
##     y8          (c)    1.271    0.132    1.003    1.525    1.009
##
## Regressions:
##                    Post.Mean  Post.SD  HPD.025  HPD.975     PSRF             Prior
```

```
##   dem60 ~
##     ind60              1.450    0.403     0.67    2.254    1.004     dnorm(0,1e-2)
##   dem65 ~
##     ind60              0.519    0.260    0.038    0.986    1.013     dnorm(0,1e-2)
##     dem60              0.897    0.087    0.765    1.144    1.041     dnorm(0,1e-2)
##
## Covariances:
##                    Post.Mean  Post.SD  HPD.025  HPD.975     PSRF             Prior
##   y1 ~~
##     y5                 0.544    0.404   -0.212    1.358    1.011        dbeta(3,3)
##   y2 ~~
##     y4                 1.478    0.732    0.114    2.967    1.005        dbeta(3,3)
##     y6                 2.152    0.770    0.707    3.696    1.004        dbeta(3,3)
##   y3 ~~
##     y7                 0.743    0.643   -0.507    2.012    1.011        dbeta(3,3)
##   y4 ~~
##     y8                 0.424    0.506   -0.525     1.44    1.023        dbeta(3,3)
##   y6 ~~
##     y8                 1.355    0.602    0.188    2.544    1.010        dbeta(3,3)
##
## Intercepts:
##                    Post.Mean  Post.SD  HPD.025  HPD.975     PSRF             Prior
##     x1                 5.042    0.088    4.859    5.211    1.041     dnorm(5,1e-2)
##     x2                 4.765    0.185    4.378    5.107    1.047     dnorm(5,1e-2)
##     x3                 3.535    0.169    3.195    3.865    1.037     dnorm(5,1e-2)
##     y1                 5.413    0.311    4.784    6.033    1.014     dnorm(5,1e-2)
##     y2                 4.190    0.452    3.253    5.042    1.009     dnorm(5,1e-2)
##     y3                 6.504    0.416    5.644     7.31    1.011     dnorm(5,1e-2)
##     y4                 4.382    0.397    3.546    5.144    1.015     dnorm(5,1e-2)
##     y5                 5.087    0.318    4.416    5.691    1.021     dnorm(5,1e-2)
##     y6                 2.922    0.408      2.1     3.73    1.016     dnorm(5,1e-2)
##     y7                 6.138    0.388    5.324    6.876    1.022     dnorm(5,1e-2)
##     y8                 3.983    0.396    3.152    4.753    1.029     dnorm(5,1e-2)
##     ind60              0.000
##     dem60              0.000
##     dem65              0.000
##
## Variances:
##                    Post.Mean  Post.SD  HPD.025  HPD.975     PSRF             Prior
##     x1                 0.093    0.023    0.049    0.139    1.007  dlnorm(1,.1)[sd]
##     x2                 0.095    0.079        0     0.24    1.018  dlnorm(1,.1)[sd]
##     x3                 0.520    0.103    0.325    0.717    1.002  dlnorm(1,.1)[sd]
##     y1                 2.029    0.514    1.077    3.086    1.015  dlnorm(1,.1)[sd]
##     y2                 8.027    1.482    5.311   10.972    1.007  dlnorm(1,.1)[sd]
##     y3                 5.373    1.064    3.438    7.525    1.006  dlnorm(1,.1)[sd]
##     y4                 3.541    0.852    1.953     5.23    1.008  dlnorm(1,.1)[sd]
##     y5                 2.461    0.541    1.467     3.53    1.013  dlnorm(1,.1)[sd]
##     y6                 5.242    0.981     3.45    7.229    1.002  dlnorm(1,.1)[sd]
##     y7                 3.844    0.816     2.38    5.528    1.008  dlnorm(1,.1)[sd]
##     y8                 3.585    0.816    2.139    5.273    1.020  dlnorm(1,.1)[sd]
##     ind60              0.456    0.092    0.285    0.635    1.001  dlnorm(1,.1)[sd]
##     dem60              3.938    0.943    2.272    5.828    1.001  dlnorm(1,.1)[sd]
##     dem65              0.134    0.177        0    0.502    1.098  dlnorm(1,.1)[sd]
```

While the results look similar to those that one would see in **lavaan**, there are a few differences that require explanation. First, the top of the output includes two model evaluation measures:

a Laplace approximation of the marginal log-likelihood (Equation (10)) and the posterior predictive p-value (see Section 2.3). Second, the "Parameter Estimates" section contains many new columns. These are the posterior mean (`Post.Mean`), the posterior standard deviation (`Post.SD`), a 95% highest posterior density interval (`HPD.025` and `HPD.975`), the potential scale reduction factor for assessing chain convergence (`PSRF`; Gelman and Rubin 1992), and the prior distribution associated with each model parameter (`Prior`). Users can optionally obtain posterior medians, posterior modes, and effective sample sizes (number of posterior samples drawn, adjusted for autocorrelation). These can be obtained by supplying the logical arguments `postmedian`, `postmode`, and `neff` to `summary()`.

We can additionally obtain various Bayesian model assessment measures via `fitMeasures()`

```
fitMeasures(fit)

##       npar       logl        ppp        bic        dic      p_dic       waic     p_waic
##     39.000  -1550.350      0.544   3268.559   3174.152     36.726   3179.450     39.163
##      looic      p_loo margloglik
##   3180.016     39.446  -2542.457
```

where, as previously mentioned, the WAIC and LOO statistics are computed with the help of package **loo**. Other **lavaan** functions, including `parameterEstimates()` and `parTable()`, work as expected. Finally, we should mention that the posterior samples and additional diagnostics are included within each object of class **blavaan**, in the `runjags` object within the `external` slot.

4.2. Measurement Invariance

Verhagen and Fox (2013; see also Verhagen, Levy, Millsap, and Fox in press) describe the use of Bayesian methods for studying the measurement invariance of sets of items or scales. Briefly, measurement invariance means that the items or scales measure diverse groups of individuals in a fair manner. For example, two individuals with the same level of mathematics ability should receive the same score on a mathematics test (within random error), regardless of the individuals' backgrounds and demographic variables. In this section, we illustrate a Bayesian measurement invariance approach using the Holzinger and Swineford (1939) data. This section also illustrates how we can take a fitted **lavaan** object and pass it to **blavaan**.

We first fit three increasingly-restricted models to the Holzinger and Swineford (1939) data via classical methods. This is accomplished via the **lavaan** `cfa()` function:

```
HS.model <- ' visual  =~ x1 + x2 + x3
              textual =~ x4 + x5 + x6
              speed   =~ x7 + x8 + x9 '

fit1 <- cfa(HS.model, data = HolzingerSwineford1939, group = "school")

fit2 <- cfa(HS.model, data = HolzingerSwineford1939, group = "school",
            group.equal = "loadings")

fit3 <- cfa(HS.model, data = HolzingerSwineford1939, group = "school",
            group.equal = c("loadings", "intercepts"))
```

The resulting objects each include a "parameter table," which contains all model details necessary for **lavaan** estimation. To fit these models via Bayesian methods, we can merely pass the classical model's parameter table and data to **blavaan**. This is accomplished via

```
bfit1 <- bcfa(parTable(fit1), data=HolzingerSwineford1939, group="school")

bfit2 <- bcfa(parTable(fit2), data=HolzingerSwineford1939, group="school")

bfit3 <- bcfa(parTable(fit3), data=HolzingerSwineford1939, group="school")
```

We could also supply the `group.equal` argument directly to the `bcfa()` calls.

Following model estimation, we can immediately compare the models via `fitMeasures()`. For example, focusing on the first two models, we obtain

```
fitMeasures(bfit1)

##       npar       logl        ppp        bic        dic      p_dic       waic     p_waic
##     60.000  -3684.313      0.000   7710.653   7482.620     56.997   7491.606     63.740
##      looic      p_loo margloglik
##   7491.956     63.915  -3938.082

fitMeasures(bfit2)

##       npar       logl        ppp        bic        dic      p_dic       waic     p_waic
##     54.000  -3687.303      0.001   7682.431   7479.265     52.329   7486.469     57.511
##      looic      p_loo margloglik
##   7487.175     57.864  -3917.858
```

As judged by the posterior predictive p-value (`ppp`), neither of the two models provide a good fit to the data. In practice, we might stop our measurement invariance study there. However, for the purpose of demonstration, we also examine the information criteria. Based on these, we would prefer the second model on the basis of its smaller DIC, WAIC, and LOOIC values.

5. Discussion

While **blavaan** currently has the ability to estimate classical, multivariate normal SEMs, there remain many additional features to be added. We discuss some of these below, including new types of models and new features for existing models.

5.1. Models

The most glaring current omission from **blavaan** is its lack of support for discrete variables. Researchers often collect Likert data and other ordinal variables, and they want to incorporate these data into SEMs. Support for these types of variables is the highest priority in future versions. From a JAGS standpoint, these models are relatively easy: as we move from continuous to discrete data, there are minimal changes to the JAGS syntax.

From a **blavaan** standpoint, the major difficulty with discrete variables involves the fit measures. Under multivariate normality, both the conditional distributions (conditioned on the latent variables) and the marginal distributions of the observed variables remain normal.

When we have discrete data, only the conditional distributions have a known form (typically multinomial). To compute fit measures, we need to obtain the marginal distributions via numerical integration. This procedure is slow because we have to repeatedly compute the marginal distribution across thousands of parameter samples. We are currently researching optimal solutions to this issue.

In addition to discrete data, we plan to support mixture SEMs, multilevel SEMs, and latent variable interactions in future releases. JAGS code to estimate specific instances of these models is available (e.g., Cho, Preacher, and Bottge 2015; Merkle and Wang 2015), but work is needed to generally implement these features. We will need to introduce new **lavaan** model specification syntax and also deal with new JAGS issues related to sampling efficiency. We expand on the latter topic in the next section.

5.2. Features

While JAGS converges for the classical models illustrated in this paper, the samples can sometimes exhibit high autocorrelation (the autocorrelation-adjusted sample size can be obtained from the model `summary()` by using the argument `neff=TRUE`). This is especially true for models that contain a large number of latent variables (say, more than 10). High autocorrelation implies that, to obtain precise posterior estimates, we may need to draw a prohibitively-large number of posterior samples in order to obtain valid parameter summaries. We are researching ways to reduce the autocorrelation; preliminary study implies that it is sometimes helpful to sample latent variables from multivariate (vs univariate) normal distributions, but this is not possible for all models (because, e.g., some pairs of latent variables may have free covariances while others have fixed zero covariances). We need to build rules in **blavaan** for deciding when we can sample from multivariate normals.

Along with the use of multivariate normals, we also plan to add an option for MCMC sampling via Stan. Stan uses information from the model gradient in order to sample more efficiently from the posterior distribution. Further work is needed to translate **lavaan** models to Stan syntax and to see whether the same parameter expansion strategy is useful there.

In summary, **blavaan** is currently useful for estimating many types of multivariate normal SEMs, and the JAGS export feature allows researchers to extend the models in any fashion desired. As additional features are added, we hope that the package keeps pace with **lavaan** as an open set of tools for SEM estimation and study.

6. Computational Details

All results were obtained using the R system for statistical computing (R Core Team 2015) version 3.2.2, employing the add-on packages blavaan 0.1-1 and lavaan 0.5-20 (Rosseel 2012).

7. Acknowledgments

This research was partially supported by NSF grants SES-1061334 and 1460719. The authors thank Ross Jacobucci for comments that helped to improve the package.

References

Arbuckle JL (2013). *IBM SPSS Amos 22 User's Guide*. Amos Development Corporation.

Barnard J, McCulloch R, Meng XL (2000). "Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage." *Statistica Sinica*, **10**, 1281–1311.

Beaujean AA (2014). *Latent variable modeling using R*. New York: Routledge.

Bollen KA (1989). *Structural equations with latent variables*. New York: Wiley.

Chib S, Greenberg E (1998). "Analysis of multivariate probit models." *Biometrika*, **85**, 347–361.

Cho SJ, Preacher KJ, Bottge BA (2015). "Detecting intervention effects in a cluster-randomized design using multilevel structural equation modeling for binary responses." *Applied Psychological Measurement*, **39**, 627–642.

Cole DA, Ciesla JA, Steiger JH (2007). "The insidious effects of failing to include design-driven correlated residuals in latent-variable covariance structure analysis." *Psychological Methods*, **12**, 381–398.

Denwood MJ (in press). "runjags: An R package providing interface utilities, model templates, parallel computing methods and additional distributions for MCMC models in JAGS." *Journal of Statistical Software*. URL http://runjags.sourceforge.net.

Gelfand AE (1996). "Model determination using sampling-based methods." In WR Gilks, S Richardson, DJ Spiegelhalter (eds.), *Markov chain Monte Carlo in Practice*, pp. 145–162. London: Chapman & Hall.

Gelman A (2004). "Parameterization and Bayesian modeling." *Journal of the American Statistical Association*, **99**, 537–545.

Gelman A (2006). "Prior distributions for variance parameters in hierarchical models." *Bayesian Analysis*, **3**, 515–534.

Gelman A, Carlin JB, Stern HS, Rubin DB (2004). *Bayesian Data Analysis*. 2nd edition. London: Chapman & Hall.

Gelman A, Rubin DB (1992). "Inference from iterative simulation using multiple sequences (with discussion)." *Statistical Science*, **7**, 457–511.

Ghosh J, Dunson DB (2009). "Default prior distributions and efficient posterior computation in Bayesian factor analysis." *Journal of Computational and Graphical Statistics*, **18**, 306–320.

Holzinger KJ, Swineford FA (1939). *A study of factor analysis: The stability of a bi-factor solution*. Number 48 in Supplementary Educational Monograph. Chicago: University of Chicago Press.

Kass RE, Raftery AE (1995). "Bayes factors." *Journal of the American Statistical Association*, **90**, 773–795.

Lee SY (2007). *Structural equation modeling: A Bayesian approach*. Chichester: Wiley.

Lewis SM, Raftery AE (1997). "Estimating Bayes factors via posterior simulation with the Laplace-Metropolis estimator." *Journal of the American Statistical Association*, **92**, 648–655.

Lunn D, Jackson C, Best N, Thomas A, Spiegelhalter D (2012). *The BUGS book: A practical introduction to Bayesian analysis*. Chapman & Hall/CRC.

Merkle EC, Wang T (2015). "Bayesian latent variable models for the analysis of experimental psychology data."

Millar RB (2009). "Comparison of hierarchical Bayesian models for overdispersed count data using DIC and Bayes' factors." *Biometrics*, **65**, 962–969.

Muthén B, Asparouhov T (2012). "Bayesian structural equation modeling: A more flexible representation of substantive theory." *Psychological Methods*, **17**, 313–335.

Muthén LK, Muthén B (1998–2012). *Mplus user's guide*. 7th edition. Los Angeles, CA: Muthén & Muthén.

Palomo J, Dunson DB, Bollen K (2007). "Bayesian structural equation modeling." In SY Lee (ed.), *Handbook of Latent Variable and Related Models*, pp. 163–188. Elsevier.

Plummer M (2003). "JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling." In K Hornik, F Leisch, A Zeileis (eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Raftery AE (1993). "Bayesian model selection in structural equation models." In KA Bollen, JS Long (eds.), *Testing Structural Equation Models*, pp. 163–180. Beverly Hills, CA: Sage.

Rosseel Y (2012). "lavaan: An R Package for Structural Equation Modeling." *Journal of Statistical Software*, **48**(2), 1–36. URL http://www.jstatsoft.org/v48/i02/.

Song XY, Lee SY (2001). "Bayesian estimation and test for factor analysis model with continuous and polytomous data in several populations." *British Journal of Mathematical and Statistical Psychology*, **54**, 237–263.

Song XY, Lee SY (2012). *Basic and advanced Bayesian structural equation modeling: With applications in the medical and behavioral sciences*. Chichester, England: Wiley.

Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A (2002). "Bayesian measures of model complexity and fit." *Journal of the Royal Statistical Society Series B*, **64**, 583–639.

Stan Development Team (2014). *Stan Modeling Language Users Guide and Reference Manual, Version 2.5.0*. URL http://mc-stan.org/.

Vehtari A, Gelman A, Gabry J (2015a). "Efficient implementation of leave-one-out cross-validation and WAIC for evaluating fitted Bayesian models." Retrieved from http://arxiv.org/abs/1507.04544.

Vehtari A, Gelman A, Gabry J (2015b). *loo: Efficient leave-one-out cross-validation and WAIC for Bayesian models*. URL https://github.com/jgabry/loo.

Verhagen AJ, Fox JP (2013). "Bayesian tests of measurement invariance." *British Journal of Mathematical and Statistical Psychology*, **66**, 383–401.

Verhagen J, Levy R, Millsap RE, Fox JP (in press). "Evaluating evidence for invariant items: A Bayes factor applied to testing measurement invariance in IRT models." *Journal of Mathematical Psychology*.

Watanabe S (2010). "Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory." *Journal of Machine Learning Research*, **11**, 3571–3594.

Wicherts JM, Dolan CV, Hessen DJ (2005). "Stereotype Threat and Group Differences in Test Performance: A Question of Measurement Invariance." *Journal of Personality and Social Psychology*, **89**(5), 696–716.

Affiliation:

Edgar Merkle
Department of Psychological Sciences
University of Missouri
28A McAlester Hall
Columbia, MO, USA 65211
E-mail: merklee@missouri.edu
URL: http://faculty.missouri.edu/~merklee/